

1 Experiment overview

1.1 Introduction

The Large Hadron Collider (LHC) is a proton-proton collider with 14 TeV centre of mass energy and design luminosity of $10^{34} \text{ cm}^{-2} \text{ s}^{-1}$. Beam crossings are 25 ns apart and at design luminosity there are 23 interactions per crossing [1-1].

The ATLAS experiment has now entered the construction phase for many of its detector components, with a strict schedule to meet the first collisions at LHC in summer 2005. The detector concept and its physics potential have been presented in the Technical Proposal [1-2] about four years ago. Over the last two years detailed descriptions of the detector systems and their performance have been presented in the various Technical Design Reports (TDR) [1-3 to 1-13]; the complex task of their integration into the overall ATLAS detector has been recently described in the Technical Coordination TDR [1-15]. The purpose of this Detector and Physics Performance TDR is to document the expected overall performance of ATLAS as well as its rich physics potential at the LHC. The main features of the detector are briefly recalled in this chapter, however the real thrust of the present document is on performance and physics.

1.1.1 Nomenclature

The beam direction defines the z -axis, and the x - y plane is the plane transverse to the beam direction. The positive x -axis is defined as pointing from the interaction point to the centre of the LHC ring, and the positive y -axis is pointing upwards. The azimuthal angle ϕ is measured around the beam axis, and the polar angle θ is the angle from the beam axis. The pseudorapidity is defined as $\eta = -\ln \tan(\theta/2)$. The transverse momentum p_T and the transverse energy E_T , as well as the missing transverse energy E_T^{miss} and other transverse variables, are defined in the x - y plane unless stated otherwise. The distance ΔR in the pseudorapidity-azimuthal angle space is defined as $\Delta R = \sqrt{\Delta^2 \eta + \Delta^2 \phi}$.

Trajectories of charged particles can be described by five helix parameters in an ideal uniform magnetic field. The following helix parametrisation is used in ATLAS, with all quantities measured at the point of closest approach to the nominal beam axis $x = 0$, $y = 0$. Parameters in x - y plane are:

$1/p_T$	Reciprocal of the transverse momentum with respect to the beam-axis.
ϕ	Azimuthal angle, where $\tan \phi \equiv p_y/p_x$.
d_0	Transverse impact parameter, defined as the transverse distance to the beam axis at the point of closest approach; signed according to the reconstructed angular momentum of the track about the axis.

Parameters in the R - z plane are:

$\cot \theta$	Cotangent of the polar angle, where $\cot \theta \equiv p_z/p_T$;
z_0	Longitudinal impact parameter, defined as the z position of the track at the point of closest approach.

The side A of the detector is defined as the side with positive z , and the side C is the side with $z < 0$. Side B is the plane with $z = 0$.

1.2 Overall detector concept

A broad spectrum of detailed physics studies led to the overall detector concept presented in the ATLAS Technical Proposal [1-2]. The basic design criteria of the detector include the following.

- Very good electromagnetic calorimetry for electron and photon identification and measurements, complemented by full-coverage hadronic calorimetry for accurate jet and missing transverse energy (E_T^{miss}) measurements;
- High-precision muon momentum measurements, with the capability to guarantee accurate measurements at the highest luminosity using the external muon spectrometer alone;
- Efficient tracking at high luminosity for high- p_T lepton-momentum measurements, electron and photon identification, τ -lepton and heavy-flavour identification, and full event reconstruction capability at lower luminosity;
- Large acceptance in pseudorapidity (η) with almost full azimuthal angle (ϕ) coverage everywhere. The azimuthal angle is measured around the beam axis, whereas pseudorapidity relates to the polar angle (θ) where θ is the angle from the z direction.
- Triggering and measurements of particles at low- p_T thresholds, providing high efficiencies for most physics processes of interest at LHC.

The overall detector layout is shown in Figure 1-i. The magnet configuration is based on an inner thin superconducting solenoid surrounding the inner detector cavity, and large superconducting air-core toroids consisting of independent coils arranged with an eight-fold symmetry outside the calorimeters.

The Inner Detector (ID) is contained within a cylinder of length 7 m and a radius of 1.15 m, in a solenoidal magnetic field of 2 T. Pattern recognition, momentum and vertex measurements, and electron identification are achieved with a combination of discrete high-resolution semiconductor pixel and strip detectors in the inner part of the tracking volume, and continuous straw-tube tracking detectors with transition radiation capability in its outer part.

Highly granular liquid-argon (LAr) electromagnetic (EM) sampling calorimetry, with excellent performance in terms of energy and position resolution, covers the pseudorapidity range $|\eta| < 3.2$. In the end-caps, the LAr technology is also used for the hadronic calorimeters, which share the cryostats with the EM end-caps. The same cryostats also house the special LAr forward calorimeters which extend the pseudorapidity coverage to $|\eta| = 4.9$. The bulk of the hadronic calorimetry is provided by a novel scintillator-tile calorimeter, which is separated into a large barrel and two smaller extended barrel cylinders, one on each side of the barrel. The overall calorimeter system provides the very good jet and E_T^{miss} performance of the detector.

The LAr calorimetry is contained in a cylinder with an outer radius of 2.25 m and extends longitudinally to ± 6.65 m along the beam axis. The outer radius of the scintillator-tile calorimeter is 4.25 m and its half length is 6.10 m. The total weight of the calorimeter system, including the solenoid flux-return iron yoke which is integrated into the tile calorimeter support structure, is about 4 000 Tons.

The calorimeter is surrounded by the muon spectrometer. The air-core toroid system, with a long barrel and two inserted end-cap magnets, generates a large magnetic field volume with strong bending power within a light and open structure. Multiple-scattering effects are thereby minimised, and excellent muon momentum resolution is achieved with three stations of high-precision tracking chambers. The muon instrumentation also includes as a key component trigger chambers with very fast time response.

The muon spectrometer defines the overall dimensions of the ATLAS detector. The outer chambers of the barrel are at a radius of about 11 m. The half-length of the barrel toroid coils is 12.5 m, and the third layer of the forward muon chambers, mounted on the cavern wall, is located about 23 m from the interaction point. The overall weight of the ATLAS detector is about 7 000 Tons.

The primary goal of the experiment is to operate at high luminosity ($10^{34} \text{ cm}^{-2}\text{s}^{-1}$) with a detector that provides as many signatures as possible. The variety of signatures is considered to be important in the harsh environment of the LHC in order to achieve robust and redundant physics measurements with the ability of internal cross-check. The measurement of the luminosity itself will be a challenge. Precision measurements employing the total and elastic cross-sections require specialised detectors. A measurement with a precision of 5% to 10% may be obtained from the machine parameters. Alternative methods involving the production of electron and muon pairs are being considered in Chapter 13.

1.3 Magnet system

The ATLAS superconducting magnet system [1-3] can be seen in Figure 1-i, and the main parameters of its components are listed in Table 1-1. It is an arrangement of a central solenoid (CS) [1-4] providing the Inner Detector with magnetic field, surrounded by a system of three large air-core toroids generating the magnetic field for the muon spectrometer. The overall dimensions of the magnet system are 26 m in length and 20 m in diameter. The two end-cap toroids (ECT) [1-5] are inserted in the barrel toroid (BT) [1-6] at each end and line up with the CS. They have a length of 5 m, an outer diameter of 10.7 m and an inner bore of 1.65 m. The CS extends over a length of 5.3 m and has a bore of 2.4 m. The unusual configuration and large size make the magnet system a considerable challenge requiring careful engineering.

The CS provides a central field of 2 T with a peak magnetic field of 2.6 T at the superconductor itself. The peak magnetic fields on the superconductors in the BT and ECT are 3.9 and 4.1 T respectively. The performance in terms of bending power is characterised by the field integral $\int B dl$, where B is the azimuthal field component and the integral is taken on a straight line trajectory between the inner and outer radius of the toroids. The BT provides 2 to 6 Tm and the ECT contributes with 4 to 8 Tm in the 0.0-1.3 and 1.6-2.7 pseudorapidity ranges respectively. The bending power is lower in the transition regions where the two magnets overlap ($1.3 < |\eta| < 1.6$).

The position of the CS in front of the EM calorimeter demands a careful minimisation of the material in order to achieve the desired calorimeter performance. As a consequence, the CS and the LAr calorimeter share one common vacuum vessel, thereby eliminating two vacuum walls.

Table 1-1 Main parameters of the toroidal and solenoid magnets in the ATLAS magnet system

Property	Unit	Barrel Toroid	End-Cap Toroid (one)	Central Solenoid
Overall dimensions:				
Inner diameter	m	9.4	1.65	2.44
Outer diameter	m	20.1	10.7	2.63
Axial length	m	25.3	5	5.3
Number of coils	-	8	8	1
Weight:				
Conductor	Tons	118	20.5	3.8
Cold mass	Tons	370	160	5.4
Total assembly	Tons	830	239	5.7
Coils:				
Number of turns per coil	-	120	116	1173
Operating current	kA	20.5	20	7.6
Stored energy	MJ	1080	206	38
Peak field	T	3.9	4.1	2.6
Conductor (NbTi/Cu/Aluminium):				
Overall size (width x height)	mm ²	57 x 12	41 x 12	30 x 4.25
Ratio Al : Cu : NbTi	-	28 : 1.3 : 1	19 : 1.3 : 1	15.6 : 0.9 : 1
Number of strands in Rutherford cable	-	33	40	12
Strand diameter	mm	1.3	1.3	1.22
Critical current at 5T, 4.2 K	kA	58	60	20.4
Total length	km	56	12.8	9.1
Cooling requirements:				
at 4.5 K	W	1130	260	100
at 60-80 K	kW	7.0	2.4	0.50

The CS coil is designed to be as thin as possible without sacrificing the operational safety and reliability. Minimum coil material and an adequate safety margin for operation are obtained by distributing the stress uniformly between the coil components, while keeping the maximum strain due to the magnetic forces below 0.1% in the principal stress components.

Each of the three toroids consists of eight coils assembled radially and symmetrically around the beam axis. The ECT coil system is rotated by 22.5° with respect to the BT coil system in order to provide radial overlap and to optimise the bending power in the interface regions of both coil systems. The BT coils are of a flat racetrack type with two double-pancake windings made of

20.5 kA aluminium-stabilised NbTi superconductor. The windings are housed in an aluminium alloy casing. The magnetic forces are transferred to the warm structure. The coils are housed in individual cryostats taking up the forces between the coils. The toroidal structure consists of eight cryostats and the linking elements between them, called *voussoirs* and struts, that provide mechanical stability. Services are brought to the coils through a cryogenic ring linking the eight cryostats to a separate service cryostat, which provides connections to the power supply, the helium refrigerator, the vacuum systems and the control system.

Each ECT also consists of eight racetrack, double-pancake coils in an aluminium alloy housing. They are cold-linked and assembled as a single cold mass, housed in one large cryostat. Therefore the internal forces in the toroids are taken by the cold supporting structure between the coils, a different design solution than in the BT. Due to the magnetic forces, the ECT magnets are pulled into the BT and the corresponding axial forces are transferred to the BT cryostats via axial transfer points linking both magnet systems. The ECT cryostats have a classical turret for services. The cryostats rest on a rail system facilitating the movement and parking of the ECT magnets for access to the detector centre.

The magnets are indirectly cooled by forced flow of helium at 4.5 K through tubes welded on the casing of the windings. The CS is cooled via a dewar coupled to the refrigerator, whereas the BT and ECT in addition have cold helium pumps to guarantee appropriate cooling by a forced helium flow. The cooling power is supplied by a central refrigeration plant located in the side cavern and the services are distributed among the four magnets.

Electrically the eight coils of the BT are connected in series, as are the 16 coils in the two ECTs. The toroid coil systems have a 21 kA power supply and are equipped with control systems for fast and slow energy dumps. The CS is energised by an 8 kA power supply. An adequate and proven quench protection system has been designed to safely dissipate the stored energies without overheating the coil windings.

The conductor used in all the coils is a composite that consists of a flat superconducting cable located in the centre of an aluminium stabiliser with rectangular cross section. For the BT and ECT, the stabiliser is made of high-purity aluminium, while in the case of the CS doped aluminium is used to provide increased mechanical strength.

1.4 Inner Detector

The layout of the Inner Detector (ID) [1-7] is shown in Figure 1-ii. It combines high-resolution detectors at the inner radii with continuous tracking elements at the outer radii, all contained in the CS which provides a nominal magnetic field of 2 T.

The momentum and vertex resolution requirements from physics call for high-precision measurements to be made with fine-granularity detectors, given the very large track density expected at the LHC. Semiconductor tracking detectors, using silicon microstrip (SCT) [1-8] and pixel [1-9] technologies offer these features. The highest granularity is achieved around the vertex region using semi-conductor pixel detectors. The total number of precision layers must be limited because of the material they introduce, and because of their high cost. Typically, three pixel layers and eight strip layers (four space points) are crossed by each track. A large number of tracking points (typically 36 per track) is provided by the straw tube tracker (TRT) [1-8], which provides continuous track-following with much less material per point and a lower cost. The combination of the two techniques gives very robust pattern recognition and high precision in

Table 1-2 Parameters of the Inner Detector. The resolutions quoted are typical values (the actual resolution in each detector depends on the impact angle).

System	Position	Area (m ²)	Resolution σ (μm)	Channels (10 ⁶)	η coverage
Pixels	1 removable barrel layer (B-layer)	0.2	$R\phi = 12, z = 66$	16	± 2.5
	2 barrel layers	1.4	$R\phi = 12, z = 66$	81	± 1.7
	5 end-cap disks on each side	0.7	$R\phi = 12, R = 77$	43	1.7-2.5
Silicon strips	4 barrel layers	34.4	$R\phi = 16, z = 580$	3.2	± 1.4
	9 end-cap wheels on each side	26.7	$R\phi = 16, R = 580$	3.0	1.4-2.5
TRT	Axial barrel straws		170 (per straw)	0.1	± 0.7
	Radial end-cap straws		170 (per straw)	0.32	0.7-2.5
	36 straws per track				

both ϕ and z coordinates. The straw hits at the outer radius contribute significantly to the momentum measurement, since the lower precision per point compared to the silicon is compensated by the large number of measurements and the higher average radius. The relative precision of the different measurements is well matched, so that no single measurement dominates the momentum resolution. This implies that the overall performance is robust. The high density of measurements in the outer part of the tracker is also valuable for the detection of photon conversions and of V^0 decays. The latter are an important element in the signature of CP violation in the B system. In addition, the electron identification capabilities of the whole experiment are enhanced by the detection of transition-radiation photons in the xenon-based gas mixture of the straw tubes.

The outer radius of the ID cavity is 115 cm, fixed by the inner dimension of the cryostat containing the LAr EM calorimeter, and the total length is 7 m, limited by the position of the end-cap calorimeters. Mechanically, the ID consists of three units: a barrel part extending over ± 80 cm, and two identical end-caps covering the rest of the cylindrical cavity. The precision tracking elements are contained within a radius of 56 cm, followed by the continuous tracking, and finally the general support and service region at the outermost radius. In order to give uniform η -coverage over the full acceptance, the final TRT wheels at high z extend inwards to a lower radius than the other TRT end-cap wheels.

In the barrel region, the high-precision detector layers are arranged on concentric cylinders around the beam axis, while the end-cap detectors are mounted on disks perpendicular to the beam axis. The pixel layers are segmented in $R\phi$ and z , while the SCT detector uses small angle (40 mrad) stereo strips to measure both coordinates, with one set of strips in each layer measuring ϕ . The barrel TRT straws are parallel to the beam direction. All the end-cap tracking elements are located in planes perpendicular to the beam axis. The strip detectors have one set of strips running radially and a set of stereo strips at an angle of 40 mrad. The continuous tracking consists of radial straws arranged into wheels.

The basic layout parameters and the expected measurement resolutions are summarised in Table 1-2. The layout provides full tracking coverage over $|\eta| \leq 2.5$, including impact parameter measurements and vertexing for heavy-flavour and τ tagging. The secondary vertex measurement performance is enhanced by the innermost layer of pixels, at a radius of about 4 cm, as close as is practical to the beam pipe. The lifetime of such a detector will be limited by radiation damage, and may need replacement after a few years, the exact time depending on the luminosity profile. A large amount of interesting physics can be done with this detector during the initial lower-luminosity running, especially in the B sector, but physics studies have demonstrated the value of good b -tagging performance during all phases of the LHC operation, for example in the case of Higgs and supersymmetry searches. It is therefore considered very important that this innermost pixel layer (or B -layer) can be replaced to maintain the highest possible performance throughout the experiment's lifetime. The mechanical design of the pixel system allows the possibility of replacing the B -layer.

1.4.1 Pixel detector

The pixel detector [1-9] is designed to provide a very high-granularity, high-precision set of measurements as close to the interaction point as possible. The system provides three precision measurements over the full acceptance, and mostly determines the impact parameter resolution and the ability of the Inner Detector to find short-lived particles such as B hadrons and τ leptons. The two-dimensional segmentation of the sensors gives space points without any of the ambiguities associated with crossed strip geometries, but requires the use of advanced electronic techniques and interconnections for the readout. The readout chips are of large area, with individual circuits for each pixel element, including buffering to store the data while awaiting the level-1 trigger decision. Each chip must be bump-bonded to the detector substrate in order to achieve the required density of connections. In addition, the chips must be radiation hardened to withstand over 300 kGy of ionising radiation and over 5×10^{14} neutrons per cm^2 over ten years of operation. The system contains a total of 140 million detector elements, each 50 μm in the $R\phi$ direction and 300 μm in z , which are invaluable for the task of pattern recognition in the crowded environment of the LHC.

The system consists of three barrels at average radii of ~ 4 cm, 10 cm, and 13 cm, and five disks on each side, between radii of 11 and 20 cm, which complete the angular coverage. The system is designed to be highly modular, containing approximately 1 500 barrel modules and 700 disk modules, and uses only one type of support structure in the barrel and two types in the disks.

The pixel modules are designed to be identical in the barrel and the disks. Each module is 62.4 mm long and 21.4 mm wide, with 61 440 pixel elements read out by 16 chips, each serving an array of 24 by 160 pixels. The output signals are routed on the sensor surface to a hybrid on top of the chips, and from there to a separate clock-and-control integrated circuit. The modules are overlapped on the support structure in order to give hermetic coverage. The thickness of each layer is expected to be about 1.7% of a radiation length at normal incidence.

1.4.2 Semiconductor tracker

The SCT system [1-8] is designed to provide eight precision measurements per track in the intermediate radial range, contributing to the measurement of momentum, impact parameter and vertex position, as well as providing good pattern recognition by the use of high granularity.

The system is an order of magnitude larger in surface area than previous generations of silicon microstrip detectors, and in addition must face radiation levels which will alter the fundamental characteristics of the silicon wafers themselves.

The barrel SCT uses eight layers of silicon microstrip detectors to provide precision points in the $R\phi$ and z coordinates, using small angle stereo to obtain the z measurement. Each silicon detector is 6.36×6.40 cm² with 768 readout strips of 80 μ m pitch. Each module consists of four single-sided p-on-n silicon detectors. On each side of the module, two detectors are wire-bonded together to form 12.8 cm long strips. Two such detector pairs are then glued together back-to-back at a 40 mrad angle, separated by a heat transport plate, and the electronics is mounted above the detectors on a hybrid. The readout chain consists of a front-end amplifier and discriminator, followed by a binary pipeline which stores the hits above threshold until the level-1 trigger decision. The end-cap modules are very similar in construction but use tapered strips, with one set aligned radially. To obtain optimal η -coverage across all end-cap wheels, end-cap modules consist of strips of either ~ 12 cm length (at the outer radii) or 6-7 cm length (at the innermost radius).

The detector contains 61 m² of silicon detectors, with 6.2 million readout channels. The spatial resolution is 16 μ m in $R\phi$ and 580 μ m in z , per module containing one $R\phi$ and one stereo measurement. Tracks can be distinguished if separated by more than ~ 200 μ m.

The barrel modules are mounted on carbon-fibre cylinders which carry the cooling system; the four complete barrels at radii of 30.0, 37.3, 44.7 and 52.0 cm are then linked together. The end-cap modules are mounted in up to three rings onto nine wheels, which are interconnected by a space-frame. The radial range of each disk is adapted to limit the coverage to $|\eta| \leq 2.5$ by equipping each one with the minimum number of rings and by using the appropriate set of modules.

Solutions have been found to the critical issues in the system, and prototype modules have been successfully tested with beams in a magnetic field, demonstrating the required performance in resolution, signal-to-noise and speed. Modules containing both detectors and front-end electronics, irradiated to the level expected for ten years of LHC operation, have also been shown to function within specifications.

Both the pixel and the SCT systems require a very high dimensional stability, cold operation of the detectors, and the removal of the heat generated by the electronics and the detector leakage current. The structures are therefore designed with materials with as low a coefficient of thermal expansion as possible.

1.4.3 Transition radiation tracker

The TRT [1-8] is based on the use of straw detectors, which can operate at the very high rates expected at the LHC by virtue of their small diameter and the isolation of the sense wires within individual gas volumes. Electron identification capability is added by employing xenon gas to detect transition-radiation photons created in a radiator between the straws. This technique is intrinsically radiation hard, and allows a large number of measurements, typically 36, to be made on every track at modest cost. However, the detector must cope with a large occupancy and high counting rates at the LHC design luminosity.

Each straw is 4 mm in diameter and equipped with a 30 μm diameter gold-plated W-Re wire, giving a fast response and good mechanical and electrical properties for a maximum straw length of 144 cm in the barrel. The barrel contains about 50 000 straws, each divided in two at the centre, in order to reduce the occupancy, and read out at each end. The end-caps contain 320 000 radial straws, with the readout at the outer radius. The total number of electronic channels is 420 000. Each channel provides a drift-time measurement, giving a spatial resolution of 170 μm per straw, and two independent thresholds. These allow the detector to discriminate between tracking hits, which pass the lower threshold, and transition-radiation (TR) hits, which pass the higher one. The TRT is operated with a non-flammable gas mixture of 70% Xe, 20% CO_2 and 10% CF_4 , with a total volume of 3 m^3 .

The barrel section is built of individual modules with between 329 and 793 axial straws each, covering the radial range from 56 to 107 cm. The first six radial layers are inactive over the central 80 cm of their length, in order to reduce their occupancy, while providing extra coverage of the crack between the barrel and end-cap sections.

The two end-caps each consist of 18 wheels. The 14 wheels nearest the interaction point cover the radial range from 64 to 103 cm, while the last four wheels extend to an inner radius of 48 cm in order to maintain a constant number of crossed straws over the full acceptance. To avoid an unnecessary increase in the number of crossed straws and material at medium rapidity, wheels 7 to 14 have half as many straws per cm in z as the other wheels.

A primary concern in the design of this sub-system has been to obtain good performance at high occupancy and counting rate. In the barrel, the rate of hits above the lower threshold varies with radius from 6 to 18 MHz, while in the end-caps the rate varies with z from 7 to 19 MHz. The maximum rate of hits above the higher TR-threshold is 1 MHz. Within a single drift-time bin, the occupancy is about one third of that in the entire active time window. Position accuracies of about 170 μm have been achieved in tests at average straw counting rates of about 12 MHz. At these rates, only about 70% of the straws give correct drift-time measurements because of shadowing effects, but the large number of straws per track guarantees a combined measurement accuracy of better than 50 μm at the LHC design luminosity, averaged over all straws and including a systematic error of ~ 30 μm from alignment.

A good pattern recognition performance is also assured by the continuous tracking. Within the radial space available, the straw spacing has been optimised for tracking at the expense of electron identification, which would be improved by a greater path length through the radiator material and fewer straw detectors. The distribution of the straws over the maximum possible path length enhances the pattern recognition performance, by reducing the effect of loopers and interactions which can saturate small regions of the detector. The TRT provides additional discrimination between electrons and hadrons, with *e.g.* a pion rejection factor at $p_T = 20$ GeV varying with η between 20 and 100 at 90 % electron efficiency.

1.5 Calorimeters

A view of the ATLAS calorimeters [1-10] is presented in Figure 1-iii. The calorimetry consists of an electromagnetic (EM) calorimeter covering the pseudorapidity region $|\eta| < 3.2$, a hadronic barrel calorimeter covering $|\eta| < 1.7$, hadronic end-cap calorimeters covering $1.5 < |\eta| < 3.2$, and forward calorimeters covering $3.1 < |\eta| < 4.9$.

The EM calorimeter is a lead/liquid-argon (LAr) detector with accordion geometry [1-11]. Over the pseudorapidity range $|\eta| < 1.8$, it is preceded by a presampler detector, installed immediately behind the cryostat cold wall, and used to correct for the energy lost in the material (ID, cryostats, coil) upstream of the calorimeter.

The hadronic barrel calorimeter is a cylinder divided into three sections: the central barrel and two identical extended barrels. It is based on a sampling technique with plastic scintillator plates (tiles) embedded in an iron absorber [1-12]. At larger pseudorapidities, where higher radiation resistance is needed, the intrinsically radiation-hard LAr technology is used for all the calorimeters [1-11]: the hadronic end-cap calorimeter, a copper LAr detector with parallel-plate geometry, and the forward calorimeter, a dense LAr calorimeter with rod-shaped electrodes in a tungsten matrix.

The barrel EM calorimeter is contained in a barrel cryostat, which surrounds the Inner Detector cavity. The solenoid which supplies the 2 T magnetic field to the Inner Detector is integrated into the vacuum of the barrel cryostat and is placed in front of the EM calorimeter. Two end-cap cryostats house the end-cap EM and hadronic calorimeters, as well as the integrated forward calorimeter. The barrel and extended barrel tile calorimeters support the LAr cryostats and also act as the main solenoid flux return.

The pseudorapidity coverage, granularity and longitudinal segmentation of the ATLAS calorimeters are summarised in Table 1-3.

The approximately 200 000 signals from the LAr calorimeters leave the cryostats through cold-to-warm feedthroughs located between the barrel and the extended barrel tile calorimeters, and at the back of each end-cap. The electronics up to the digitisation stage will be contained in radial boxes attached to each feedthrough and located in the vertical gaps between the barrel and extended barrel tile calorimeters.

1.5.1 Electromagnetic calorimeter

The EM calorimeter [1-11] is divided into a barrel part ($|\eta| < 1.475$) and two end-caps ($1.375 < |\eta| < 3.2$). The barrel calorimeter consists of two identical half-barrels, separated by a small gap (6 mm) at $z = 0$. Each end-cap calorimeter is mechanically divided into two coaxial wheels: an outer wheel covering the region $1.375 < |\eta| < 2.5$, and an inner wheel covering the region $2.5 < |\eta| < 3.2$.

The EM calorimeter is a lead LAr detector with accordion-shaped Kapton electrodes and lead absorber plates over its full coverage. The accordion geometry provides complete ϕ symmetry without azimuthal cracks. The lead thickness in the absorber plates has been optimised as a function of η in terms of EM calorimeter performance in energy resolution. The LAr gap has a constant thickness of 2.1 mm in the barrel. In the end-cap, the shape of the Kapton electrodes and lead converter plates is more complicated, because the amplitude of the accordion waves increases with radius. The absorbers have constant thickness, and therefore the LAr gap also increases with radius. The total thickness of the EM calorimeter is > 24 radiation lengths (X_0) in the barrel and $> 26 X_0$ in the end-caps.

Over the region devoted to precision physics ($|\eta| < 2.5$), the EM calorimeter is segmented into three longitudinal sections. The strip section, which has a constant thickness of $\sim 6 X_0$ (upstream material included) as a function of η , is equipped with narrow strips with a pitch of ~ 4 mm in the η direction. This section acts as a 'preshower' detector, enhancing particle identification

Table 1-3 Pseudorapidity coverage, granularity and longitudinal segmentation of the ATLAS calorimeters.

EM CALORIMETER	Barrel	End-cap	
Coverage	$ \eta < 1.475$	$1.375 < \eta < 3.2$	
Longitudinal segmentation	3 samplings	3 samplings	$1.5 < \eta < 2.5$
		2 samplings	$1.375 < \eta < 1.5$
			$2.5 < \eta < 3.2$
Granularity ($\Delta\eta \times \Delta\phi$)			
Sampling 1	0.003×0.1	0.025×0.1	$1.375 < \eta < 1.5$
		0.003×0.1	$1.5 < \eta < 1.8$
		0.004×0.1	$1.8 < \eta < 2.0$
		0.006×0.1	$2.0 < \eta < 2.5$
		0.1×0.1	$2.5 < \eta < 3.2$
Sampling 2	0.025×0.025	0.025×0.025	$1.375 < \eta < 2.5$
		0.1×0.1	$2.5 < \eta < 3.2$
Sampling 3	0.05×0.025	0.05×0.025	$1.5 < \eta < 2.5$
PRESAMPLER	Barrel	End-cap	
Coverage	$ \eta < 1.52$	$1.5 < \eta < 1.8$	
Longitudinal segmentation	1 sampling	1 sampling	
Granularity ($\Delta\eta \times \Delta\phi$)	0.025×0.1	0.025×0.1	
HADRONIC TILE	Barrel	Extended barrel	
Coverage	$ \eta < 1.0$	$0.8 < \eta < 1.7$	
Longitudinal segmentation	3 samplings	3 samplings	
Granularity ($\Delta\eta \times \Delta\phi$)			
Samplings 1 and 2	0.1×0.1	0.1×0.1	
Sampling 3	0.2×0.1	0.2×0.1	
HADRONIC LAr	End-cap		
Coverage	$1.5 < \eta < 3.2$		
Longitudinal segmentation	4 samplings		
Granularity ($\Delta\eta \times \Delta\phi$)		0.1×0.1	$1.5 < \eta < 2.5$
		0.2×0.2	$2.5 < \eta < 3.2$
FORWARD CALORIMETER	Forward		
Coverage	$3.1 < \eta < 4.9$		
Longitudinal segmentation	3 samplings		
Granularity ($\Delta\eta \times \Delta\phi$)	$\sim 0.2 \times 0.2$		

(γ/π^0 , e/π separation, etc.) and providing a precise position measurement in η . The middle section is transversally segmented into square towers of size $\Delta\eta \times \Delta\phi = 0.025 \times 0.025$ ($\sim 4 \times 4$ cm² at $\eta = 0$). The total calorimeter thickness up to the end of the second section is $\sim 24 X_0$, tapered with increasing rapidity (this includes also the upstream material). The back section has a granularity of 0.05 in η and a thickness varying between $2 X_0$ and $12 X_0$. For $|\eta| > 2.5$, i.e. for the end-cap inner wheel, the calorimeter is segmented in two longitudinal sections and has a coarser lateral

granularity than for the rest of the acceptance. This is sufficient to satisfy the physics requirements (reconstruction of jets and measurement of E_T^{miss}). The calorimeter cells point towards the interaction region over the complete η -coverage. The total number of channels is $\sim 190\,000$.

The total material seen by an incident particle before the calorimeter front face is about $2.3 X_0$ at $\eta = 0$, and increases with pseudorapidity in the barrel because of the particle angle. In the region where the amount of material exceeds $\sim 2 X_0$ (as is the case for $|\eta| < 1.8$), a presampler is used to correct for the energy lost by electrons and photons upstream of the calorimeter. The presampler consists of an active LAr layer of thickness 1.1 cm (0.5 cm) in the barrel (end-cap) region. At the transition between the barrel and the end-cap calorimeters, *i.e.* at the boundary between the two cryostats, the amount of material in front of the calorimeter reaches a localised maximum of about $7 X_0$. In this region, the presampler is complemented by a scintillator slab inserted in the crack between the barrel and end-cap cryostats and covering the region $1.0 < |\eta| < 1.6$. The region $1.37 < |\eta| < 1.52$ is not used for precision physics measurements involving photons because of the large amount of material situated in front of the EM calorimeter.

The signals from the EM calorimeters are extracted at the detector inner and outer faces and sent to preamplifiers located outside the cryostats close to the feedthroughs. Cables of 25Ω are used for cells belonging to the second and third samplings, while 50Ω cables are used for the strips in the first sampling and for the presampler pads, which have lower capacitance, in order to optimise the electronic noise performance.

The preamplifier output is formed by bipolar shapers, sampled every 25 ns, and stored in analogue memories using Switching Capacitor Arrays (SCA) during the level-1 trigger latency. If the level-1 trigger is validated, the corresponding samples (typically five) are extracted from the SCA, digitised and read out to the data acquisition system.

1.5.2 Hadronic calorimeters

The ATLAS hadronic calorimeters cover the range $|\eta| < 4.9$ using different techniques best suited for the widely varying requirements and radiation environment over the large η -range. Over the range $|\eta| < 1.7$, the iron scintillating-tile technique is used for the barrel and extended barrel tile calorimeters and for partially instrumenting the gap between them with the intermediate tile calorimeter (ITC). This gap provides space for cables and services from the innermost detectors. Over the range $\sim 1.5 < |\eta| < 4.9$, LAr calorimeters were chosen: the hadronic end-cap calorimeter (HEC) extends to $|\eta| < 3.2$, while the range $3.1 < |\eta| < 4.9$ is covered by the high-density forward calorimeter (FCAL). Both the HEC and the FCAL are integrated in the same cryostat as that housing the EM end-caps. Table 1-3 shows the granularity of the hadronic calorimetry over the full η -range.

An important parameter in the design of the hadronic calorimeter is its thickness: it has to provide good containment for hadronic showers and reduce punch-through into the muon system to a minimum. The total thickness is 11 interaction lengths (λ) at $\eta = 0$, including about 1.5λ from the outer support, which has been shown both by measurements and simulation to be sufficient to reduce the punch-through well below the irreducible level of prompt or decay muons. Close to 10λ of active calorimeter are adequate to provide good resolution for high energy jets. Together with the large η -coverage, this will also guarantee a good E_T^{miss} measurement, which is important for many physics signatures and in particular for SUSY particle searches.

1.5.2.1 Tile calorimeter

The large hadronic barrel calorimeter [1-12] is a sampling calorimeter using iron as the absorber and scintillating tiles as the active material. The tiles are placed radially and staggered in depth. The structure is periodic along z . The tiles are 3 mm thick and the total thickness of the iron plates in one period is 14 mm. Two sides of the scintillating tiles are read out by wavelength shifting (WLS) fibres into two separate photomultipliers (PMTs).

The tile calorimeter is composed of one barrel and two extended barrels. Radially the tile calorimeter extends from an inner radius of 2.28 m to an outer radius of 4.25 m. It is longitudinally segmented in three layers, approximately 1.4, 4.0 and 1.8 interaction lengths thick at $\eta = 0$. Azimuthally, the barrel and extended barrels are divided into 64 modules. In η , the readout cells, built by grouping fibres into PMTs, are ‘pseudo-projective’ towards the interaction region. The resulting granularity is $\Delta\eta \times \Delta\phi = 0.1 \times 0.1$ (0.2×0.1 in the last layer), as shown in Table 1-3. The total number of channels is about 10 000. The calorimeter is placed behind the EM calorimeter ($\approx 1.2 \lambda$) and the solenoid coil. The total thickness at the outer edge of the tile-instrumented region is 9.2λ at $\eta=0$.

The barrel cylinder covers the region $|\eta| < 1.0$. A vertical gap of 68 cm width provides space for cables from the ID, feedthroughs, and service pipes for the EM calorimeter and the CS; it also houses front-end electronics for the EM calorimeter. The extended barrel covers the region $0.8 < |\eta| < 1.7$. The azimuthal segmentation is as for the barrel, but the radial segmentation differs for the second and third layers. The thickness of the calorimeter in the gap is improved by the ITC, which has the same segmentation as the rest of the tile calorimeter. It is composed of two radial sections attached on the face of the extended barrel. The outer section, 31 cm thick, starts at the outer radius and covers 45 cm in radius. It is followed by the inner section which is 9 cm thick and extends over 45 cm to lower radii. The ITC is extended further inwards by a scintillator sheet, covering the inner part of the extended barrel and extending to the region between the LAr barrel and end-cap cryostats over $1.0 < |\eta| < 1.6$. This scintillator samples the energy lost in the cryostat walls and dead material. It is segmented in three sections of $\Delta\eta \sim 0.2$.

The signals produced by the scintillating tiles and collected by the WLS fibres are fast. The PMTs have low dark current and are also fast (rise time and transit time of a few ns). The shaper transforms the current pulse from the PMT into a unipolar pulse of FWHM of 50 ns.

1.5.2.2 Liquid-argon hadronic end-cap calorimeters

Each HEC [1-11] consists of two independent wheels, of outer radius 2.03 m. The upstream wheel is built out of 25 mm copper plates, while the cheaper other one, farther from the interaction point, uses 50 mm plates. In both wheels, the 8.5 mm gap between consecutive copper plates is equipped with three parallel electrodes, splitting the gap into four drift spaces of about 1.8 mm. The readout electrode is the central one, which is a three layer printed circuit, as in the EM calorimeter. The two layer printed circuits on either side serve only as high-voltage carriers. This electrode structure forms an ‘electrostatic transformer’ (EST) with an EST ratio of two. Such a scheme has the same behaviour as a double gap of 4 mm, but without the drawbacks associated with very high voltage (typically 4 kV instead of 2 kV), and ion build up in larger gaps.

Each wheel is built out of 32 identical modules, assembled with fixtures at the periphery, and a central ring. The central (buried) layer of the readout boards features a pad structure which defines the transverse readout granularity. The other layers are made out of a high resistivity coating, with a typical surface resistance of 1 M Ω per square. Each wheel is divided into two longitudinal segments. The weight of the first (second) wheel is 67 (90) Tons.

Primarily in order to limit the capacitance seen by a single preamplifier, and thus to allow for a fast response, only two gaps are ganged together at the pad level. Miniature coaxial cables running between the sectors carry signals to the preamplifier boards located at the wheel periphery. Output signals from (typically) four preamplifiers are summed together on the same board. A buffer stage drives the output signal up to the cold-to-warm feedthroughs.

Cells defined in this way are fully projective in azimuth, but only 'pseudo-projective' in η . However, the detector envelope is cylindrical, for the sake of mechanical simplicity. To minimise the dip in the material density at the transition between the end-cap and the forward calorimeter (around $|\eta| = 3.1$), the end-cap EM calorimeter reaches $|\eta| = 3.2$, thereby overlapping the forward calorimeter.

1.5.2.3 Liquid-argon forward calorimeter

The FCAL [1-11] is a particularly challenging detector owing to the high level of radiation it has to cope with. In ATLAS, the forward calorimeter is integrated into the end-cap cryostat, with a front face at about 4.7 m from the interaction point. Compared to layouts with a forward calorimeter situated at much larger distances from the interaction point, the survival of such a calorimeter in terms of radiation resistance is clearly more difficult. On the other hand, the integrated FCAL provides clear benefits in terms of uniformity of the calorimetric coverage as well as reduced radiation background levels in the muon spectrometer.

A global view of the FCAL is presented in Figure 1-iii. In order to minimise the amount of neutron albedo in the ID cavity (as an indication, at 47 cm radius and z less than 1 m, the albedo from the forward calorimeter contributes less than 5% of the total neutron fluence), the front face of the FCAL is recessed by about 1.2 m with respect to the EM calorimeter front face. This severely limits longitudinal space for installing about 9.5 active interaction lengths, and therefore calls for a high-density design, which also avoids energy leakage from the FCAL to its neighbours.

The FCAL consists of three sections: the first one is made of copper, while the other two are made out of tungsten. In each section the calorimeter consists of a metal matrix with regularly spaced longitudinal channels filled with concentric rods and tubes. The rods are at positive high voltage while the tubes and matrix are grounded. The LAr in the gap between is the sensitive medium. This geometry allows for an excellent control of the gaps which are as small as 250 μm in the first section.

While the construction of the copper section does not present special difficulties, the construction of a tungsten calorimeter is a rather new and challenging task. After successful assembly of several prototypes, a technique has been chosen based on the assembly of small sintered tungsten alloy pieces. The overall density (including the liquid argon) of a section built in this way, with 375 μm gaps, is 14.5 g/cm³.

Particular care is needed to support the FCAL in the end-cap cryostat such that the sensitive area is extended down to an angle as small as possible. An external structural tube carries the weight of the forward calorimeter, and withstands the pressure on the cryostat end-walls, while the central cryostat tube near the beam pipe has no structural role.

In terms of electronics and readout, four rods are ganged on the detector, and the signal is carried out by polyimide insulated coaxial cables. The number of channels is 3 584 for the total of both sides.

1.6 Muon spectrometer

The conceptual layout of the muon spectrometer [1-13] is visible in Figure 1-i. It is based on the magnetic deflection of muon tracks in the large superconducting air-core toroid magnets, instrumented with separate trigger and high-precision tracking chambers. Over the range $|\eta| \leq 1.0$, magnetic bending is provided by the large barrel toroid. For $1.4 \leq |\eta| \leq 2.7$, muon tracks are bent by two smaller end-cap magnets inserted into both ends of the barrel toroid. Over $1.0 \leq |\eta| \leq 1.4$, usually referred to as the transition region, magnetic deflection is provided by a combination of barrel and end-cap fields. This magnet configuration provides a field that is mostly orthogonal to the muon trajectories, while minimising the degradation of resolution due to multiple scattering.

The anticipated high level of particle fluxes has had a major impact on the choice and design of the spectrometer instrumentation, affecting required performance parameters such as rate capability, granularity, ageing properties and radiation hardness. Trigger and reconstruction algorithms have been optimised to cope with the difficult background conditions resulting from penetrating primary collision products and from radiation backgrounds, mostly neutrons and photons in the 1 MeV range, produced from secondary interactions in the calorimeters, shielding material, beam pipe and LHC machine elements.

In the barrel region, tracks are measured in chambers arranged in three cylindrical layers ('stations') around the beam axis; in the transition and end-cap regions, the chambers are installed vertically, also in three stations. Over most of the η -range, a precision measurement of the track coordinates in the principal bending direction of the magnetic field is provided by Monitored Drift Tubes (MDTs). At large pseudorapidities and close to the interaction point, Cathode Strip Chambers (CSCs) with higher granularity are used in the innermost plane over $2 < |\eta| < 2.7$, to withstand the demanding rate and background conditions. Optical alignment systems have been designed to meet the stringent requirements on the mechanical accuracy and the survey of the precision chambers.

The precision measurement of the muon tracks is made in the R - z projection, in a direction parallel to the bending direction of the magnetic field; the axial coordinate (z) is measured in the barrel and the radial coordinate (R) in the transition and end-cap regions. The MDTs provide a single-wire resolution of $\sim 80 \mu\text{m}$ when operated at high gas pressure (3 bar) together with robust and reliable operation thanks to the mechanical isolation of each sense wire from its neighbours. The construction of prototypes has demonstrated that the MDTs can be built to the required mechanical accuracy of $\sim 30 \mu\text{m}$.

The trigger system covers the pseudorapidity range $|\eta| \leq 2.4$. Resistive Plate Chambers (RPCs) are used in the barrel and Thin Gap Chambers (TGCs) in the end-cap regions. The trigger chambers for the ATLAS muon spectrometer serve a threefold purpose:

- bunch crossing identification, requiring a time resolution better than the LHC bunch spacing of 25 ns;
- a trigger with well-defined p_T cut-offs in moderate magnetic fields, requiring a granularity of the order of 1 cm;
- measurement of the second coordinate in a direction orthogonal to that measured by the precision chambers, with a typical resolution of 5–10 mm.

1.6.1 Muon chamber layout

The overall layout of the muon chambers in the ATLAS detector is shown in Figure 1-1, which indicates the different regions in which the four chamber technologies described above are employed. The chambers are arranged such that particles from the interaction point traverse three stations of chambers. The positions of these stations are optimised for essentially full coverage and momentum resolution. In the barrel, particles are measured near the inner and outer field boundaries, and inside the field volume, in order to determine the momentum from the sagitta of the trajectory. In the end-cap regions, for $|\eta| > 1.4$, the magnet cryostats do not allow the positioning of chambers inside the field volume. Instead, the chambers are arranged to determine the momentum with the best possible resolution from a point-angle measurement (this is also the case in the barrel region in the vicinity of the coils).




Figure 1-1 Three-dimensional view of the muon spectrometer instrumentation indicating the areas covered by the four different chamber technologies

The barrel chambers form three cylinders concentric with the beam axis, at radii of about 5, 7.5, and 10 m. They cover the pseudorapidity range $|\eta| < 1$. The end-cap chambers cover the range $1 < |\eta| < 2.7$ and are arranged in four disks at distances of 7, 10, 14, and 21–23 m from the inter-

action point, concentric with the beam axis. All chambers combined provide almost complete coverage of the pseudorapidity range $1.0 < |\eta| < 2.7$. There is an opening in the central R- ϕ plane ($\eta = 0$) for the passage of cables and services of the ID, the CS, and the calorimeters.

In the barrel, the chambers are arranged in projective towers. Particles are measured in 2×4 sensitive layers in the inner station and in 2×3 layers each in the middle and outer stations. Within a projective tower, the chambers are optically connected by alignment rays which monitor the relative chamber positions. A different alignment strategy is used in the end-caps, where the positions of complete chamber planes are monitored. No active repositioning of the chambers is foreseen.

Both in the barrel and the end-caps, a 16-fold segmentation in azimuth has been chosen that follows the eightfold azimuthal symmetry of the magnet structure. The chambers are arranged in large and small sectors. The large sectors cover the regions between the BT coils and the small sectors the azimuthal range around the BT coils. In two lower barrel sectors, the rails carrying the calorimeter and their feet require specially shaped chambers to maximise the detector acceptance.

The trigger function in the barrel is provided by three stations of RPCs. They are located on both sides of the middle MDT station, and directly inside the outer MDT station. In the end-caps, the trigger is provided by three stations of TGCs located near the middle MDT station.

Maximum standardisation and the smallest possible number of different chamber sizes have also been important goals of the detector layout. The barrel chambers are of rectangular shape with areas of 2–10 m². The end-cap chambers are of trapezoidal shape ('staircase' approximation) with tapering angles of 8.5° and 14° for the small and large chambers, respectively. Their areas range from 1–10 m² for individual chamber modules and up to 30 m² when several of them are preassembled for installation. Table 1-4 summarises the number of chambers, the area covered, and the number of readout channels for the four chamber technologies.

Table 1-4 Overview of the muon chamber instrumentation. 'Area covered' refers to chamber modules which normally contain several detector layers.

	Precision chambers		Trigger chambers	
	CSC	MDT	RPC	TGC
Number of chambers	32	1194	596	192
Number of readout channels	67 000	370 000	355 000	440 000
Area covered (m ²)	27	5500	3650	2900

All the barrel chambers and a part of the end-cap chambers are supported by the barrel toroid structure. Wherever possible, chambers are installed on rails parallel to the beam axis and connected to the magnet. They will be installed by sliding them in from the two detector ends. The majority of the end-cap chambers are mounted on the service structures at the two extreme ends of the experimental hall, and on special muon chamber support frames.

1.6.2 Monitored drift-tube chambers

The basic detection elements of the MDT chambers [1-13] are aluminium tubes of 30 mm diameter and 400 μm wall thickness, with a 50 μm diameter central W-Re wire. The tubes are operated with a non-flammable mixture of 93% Ar and 7% CO₂ at 3 bar absolute pressure and have a total volume of 800 m³. The chosen working point provides for a non-linear space-time relation with a maximum drift time of ~ 700 ns, a small Lorentz angle, and excellent ageing properties. The single-wire resolution is ~ 80 μm .

The tubes are produced by extrusion from a hard aluminium alloy, and are available commercially. They are closed by endplugs which provide for accurate positioning of the anode wires, wire tension, gas tightness, and electrical and gas connections. The drift tubes can be manufactured to tight mechanical tolerances which are well matched to their intrinsic resolution properties, mostly using automated assembly procedures. The tube lengths vary from 70 cm to 630 cm.

To improve the resolution of a chamber beyond the single-wire limit and to achieve adequate redundancy for pattern recognition, the MDT chambers are constructed from 2 \times 4 monolayers of drift tubes for the inner station and 2 \times 3 monolayers for the middle and outer stations. The tubes are arranged in multilayer pairs of three or four monolayers, respectively, on opposite sides of a rigid support structure. The support structures ('spacer frames') provide for accurate positioning of the drift tubes with respect to each other, and for mechanical integrity under effects of temperature and gravity; for the barrel chambers which are not mounted in a vertical plane, they are designed to bend the drift tubes slightly in order to match them to the gravitational sag of the wires. The spacer frames also support most of the components of the alignment system.

The structural components of the spacer frames are three 'cross-plates', to which the drift-tube multilayers are attached, and two 'long beams' connecting the cross-plates. The frames need to be constructed to a moderate mechanical accuracy of 0.5 mm only; accurate positioning of the drift tubes is provided by the assembly procedure. They will be attached to the rail structures of the spectrometer by three-point kinematic supports. Once a chamber is installed in its final location in the spectrometer, mechanical deformations are monitored by an in-plane optical system; hence the name 'monitored drift-tube chambers'.

Each drift tube is read out at one end by a low-impedance current sensitive preamplifier, with a threshold five times above the noise level. The preamplifier is followed by a differential amplifier, a shaping amplifier and a discriminator. The output of the shaping amplifier is also connected to a simple ADC, to correct the drift-time measurement for time-slewing using the charge-integrated signal.

1.6.3 Cathode strip chambers

The CSCs [1-13] are multiwire proportional chambers with cathode strip readout and with a symmetric cell in which the anode-cathode spacing is equal to the anode wire pitch. The precision coordinate is obtained by measuring the charge induced on the segmented cathode by the avalanche formed on the anode wire. Good spatial resolution is achieved by segmentation of the readout cathode and by charge interpolation between neighbouring strips. The cathode strips for the precision measurement are orthogonal to the anode wires. The anode wire pitch is 2.54 mm and the cathode readout pitch is 5.08 mm; position resolutions of better than 60 μm have been measured in several prototypes. Other important characteristics are small electron

drift times (30 ns), good time resolution (7 ns), good two-track resolution, and low neutron sensitivity. A measurement of the transverse coordinate is obtained from orthogonal strips, *i.e.* oriented parallel to the anode wires, which form the second cathode of the chamber.

The spatial resolution of the CSCs is sensitive to the inclination of tracks and the Lorentz angle. To minimise degradations of the resolution due to these effects, they will be installed in a tilted position such that stiff tracks originating from the interaction point are normal to the chamber surface. The CSCs are arranged in 2×4 layers.

The baseline CSC gas is a non-flammable mixture of 30% Ar, 50% CO₂ and 20% CF₄, with a total volume of 1.1 m³. The fact that this gas contains no hydrogen, combined with the small gap width, explains the low sensitivity to neutron backgrounds.

The front-end section of the strip readout electronics consists of a charge-sensitive preamplifier that drives a pulse-shaping amplifier. This is followed by analogue storage of the peak cathode pulse height during the level-1 trigger latency. After a level-1 trigger, the analogue data are multiplexed into a 10-bit ADC. Since the precision coordinate is obtained from charge interpolation, the position resolution depends critically on the relative gain of neighbouring cathode strips and readout channels, which is measured with the help of accurately calibrated test pulses.

1.6.4 Resistive plate chambers

The RPC [1-13] is a gaseous detector providing a typical space-time resolution of 1 cm \times 1 ns with digital readout. The basic RPC unit is a narrow gas gap formed by two parallel resistive bakelite plates, separated by insulating spacers. The primary ionisation electrons are multiplied into avalanches by a high, uniform electric field of typically 4.5 kV/mm. Amplification in avalanche mode produces pulses of typically 0.5 pC. The gas mixture is based on tetrafluoroethane (C₂H₂F₄) with some small admixture of SF₆, a non-flammable and environmentally safe gas that allows for a relatively low operating voltage. The signal is read out via capacitive coupling by metal strips on both sides of the detector. A trigger chamber is made from two rectangular detector layers, each one read out by two orthogonal series of pick-up strips: the ‘ η strips’ are parallel to the MDT wires and provide the bending view of the trigger detector; the ‘ ϕ strips’, orthogonal to the MDT wires, provide the second-coordinate measurement which is also required for the offline pattern recognition.

The RPCs use no wires and therefore have a simple mechanical structure and are straightforward to manufacture. The 2 mm thick bakelite plates are separated by polycarbonate spacers of 2 mm thickness which define the size of the gas gap. The spacers are glued on both plates at 10 cm intervals. A 7 mm wide frame of the same material and thickness as the spacers is used to seal the gas gap at all four edges. The outside surfaces of the resistive plates are coated with thin layers of graphite paint which are connected to the high voltage supply. These graphite electrodes are separated from the pick-up strips by 200 μ m thick insulating films which are glued on both graphite layers. The readout strips are arranged with a pitch varying from 30.0 to 39.5 mm.

Each chamber is made from two detector layers and four readout strip panels. These elements are rigidly held together by two support panels which provide the required mechanical stiffness of the chambers. The panels are made of polystyrene sandwiched between two aluminium sheets. One panel is flat, 50 mm thick, with 0.5 mm thick aluminium sheets; the other panel is

10 mm thick with 0.3 mm sheets and is preloaded with a 1 cm sagitta. The two panels are rigidly connected by 2 mm thick aluminium profiles, such that the preloaded support panel provides uniform pressure over the whole surface of an RPC module.

To preserve the excellent intrinsic time resolution of the RPCs, the readout strips are optimised for good transmission properties and are terminated at both ends to avoid signal reflections. The front-end electronics are based on a three-stage voltage amplifier followed by a variable-threshold comparator. They are mounted on printed circuit boards attached to the edges of the readout panels.

1.6.5 Thin gap chambers

The TGCs [1-13] are similar in design to multiwire proportional chambers, with the difference that the anode wire pitch is larger than the cathode-anode distance. Signals from the anode wires, arranged parallel to the MDT wires, provide the trigger information together with readout strips arranged orthogonal to the wires. These readout strips are also used to measure the second coordinate.

Operated with a highly quenching gas mixture of 55% CO₂ and 45% *n*-pentane (*n*-C₅H₁₂), this type of cell geometry permits operation in saturated mode, with a number of advantages:

- small sensitivity to mechanical deformations, which is important to minimise the cost of large-area chambers;
- small dependence of the pulse height on the incident angle, up to angles of 40°;
- nearly Gaussian pulse height distribution with small Landau tails, and no streamer formation.

The total gas volume is 16 m³. The gas mixture is highly flammable and requires adequate safety precautions.

The main dimensional characteristics of the chambers are a cathode-cathode distance (gas gap) of 2.8 mm, a wire pitch of 1.8 mm, and a wire diameter of 50 μm. The operating high voltage foreseen is 3.1 kV. The electric field configuration and the small wire distance provide for a short drift time and thus a good time resolution. Ageing properties of the chambers have been investigated in detail and were found to be fully adequate for the expected operating conditions at the LHC, including the large safety margin used throughout the muon spectrometer chambers to account for uncertainties in the predicted occupancies from the neutron background.

The TGCs are constructed in doublets and in triplets of chambers. The inner station consists of one doublet and is only used to measure the second coordinate. The seven chamber layers in the middle station are arranged in one triplet and two doublets which provide the trigger and the second coordinate measurements. The anode plane is sandwiched between two cathode planes made of 1.6 mm G-10 plates on which the graphite cathode is deposited. On the backside of the cathode plates facing the centre plane of the chamber, etched copper strips provide the readout of the azimuthal coordinate; no readout strips are foreseen for the central layer of a triplet. The TGC layers are separated by 20 mm thick paper honeycomb panels which provide a rigid mechanical structure for the chambers. On the outside, the gas pressure is sustained by 5 mm thick paper honeycomb panels. These are covered in turn by 0.5 mm G-10 plates.

To form a trigger signal, several anode wires are grouped together and fed to a common readout channel. The number of wires per group varies between 4 and 20, depending on the desired granularity as a function of pseudorapidity. The ganged signals are fed into a low-impedance two-stage amplifier.

1.6.6 Alignment

The requirements on the momentum resolution of the spectrometer call for an accuracy of the relative positioning of chambers traversed by a muon track that matches the intrinsic resolution and the mechanical tolerances of the precision chambers. Over the large global dimensions of the spectrometer, however, it is not possible to stabilise the dimensions and positions of the chambers at the $30\ \mu\text{m}$ level. Therefore, chamber deformations and positions are constantly monitored by means of optical alignment systems and displacements up to $\sim 1\ \text{cm}$ can readily be corrected for in the offline analysis.

All alignment systems are based on optically monitoring deviations from straight lines. Owing to geometrical constraints, different schemes are used to monitor chamber positions in the barrel, in the end-caps, and the deformations of large chambers ('in-plane alignment'). For reasons of cost, optical monitoring in the barrel is foreseen only for the large sectors of chambers. Chambers in the small sectors are aligned with particle tracks, exploiting the overlap with chambers in the large sectors. Alignment with tracks will also serve to cross-calibrate the optical survey of the large sectors.

The very high accuracy of $30\ \mu\text{m}$ is required only for the positioning of chambers within a projective tower. The accuracy required for the relative positioning of different towers to obtain adequate mass resolutions for multimMuon final states is in the millimetre range. This accuracy is easily achieved by the initial positioning and survey of chambers at installation time. The relative alignment of muon spectrometer, calorimeters and Inner Detector will rely on high-momentum muon trajectories.

1.7 Trigger and data-acquisition system

The ATLAS trigger and data-acquisition (DAQ) system is based on three levels of online event selection [1-14],[1-16],[1-17]. Each trigger level refines the decisions made at the previous level and, where necessary, applies additional selection criteria. Starting from an initial bunch-crossing rate of 40 MHz (interaction rate of $\sim 10^9\ \text{Hz}$ at a luminosity of $10^{34}\ \text{cm}^{-2}\text{s}^{-1}$), the rate of selected events must be reduced to $\sim 100\ \text{Hz}$ for permanent storage. While this requires an overall rejection factor of 10^7 against 'minimum-bias' events, excellent efficiency must be retained for the rare new physics processes, such as Higgs boson decays, which will be searched for in ATLAS. Figure 1-2 shows a simplified functional view of the Trigger/DAQ system. In the following, a brief description is given of some of the key aspects of the event-selection process.

The level-1 (LVL1) trigger [1-14] makes an initial selection based on reduced-granularity information from a subset of detectors. High transverse-momentum (high- p_T) muons are identified using only the trigger chambers, RPCs in the barrel, and TGCs in the end-caps. The calorimeter selections are based on reduced-granularity information from all the calorimeters (EM and hadronic; barrel, end-cap and forward). Objects searched for by the calorimeter trigger are high- p_T electrons and photons, jets, and τ -leptons decaying into hadrons, as well as large missing and total transverse energies. In the case of the electron/photon and hadron/ τ triggers, energy iso-




Figure 1-2 Block diagram of the Trigger/DAQ system.

lation cuts can be applied. Trigger information is provided for a number of sets of p_T thresholds (generally 6–8 sets of thresholds per object type). The missing and total scalar transverse energies used in the LVL1 trigger are calculated by summing over trigger towers. In addition, a trigger on the scalar sum of jet transverse energies is also available.

The LVL1 trigger decision is based on combinations of objects required in coincidence or veto. Most of the physics requirements of ATLAS can be met by using, at the LVL1 trigger level, fairly simple selection criteria of a rather inclusive nature. However, the trigger implementation is flexible and it can be programmed to select events using more complicated signatures.

The maximum rate at which the ATLAS front-end systems can accept LVL1 triggers is limited to 75 kHz (upgradable to 100 kHz). The rates estimated in trigger performance studies, using trigger menus that meet the needs of the ATLAS physics programme, are about a factor of two below this limit. Given that there are large intrinsic uncertainties in the calculations, this safety factor is not over-generous. However, if necessary, rates could be significantly reduced without major consequences for the physics programme, for example by increasing the thresholds on some of the inclusive (single-object) triggers when operating at the highest luminosities, and by relying more heavily on multi-object triggers.

An essential requirement on the LVL1 trigger is that it should uniquely identify the bunch-crossing of interest. Given the short (25 ns) bunch-crossing interval, this is a non-trivial consideration. In the case of the muon trigger, the physical size of the muon spectrometer implies times-of-flight comparable to the bunch-crossing period. For the calorimeter trigger, a serious challenge is that the pulse shape of the calorimeter signals extends over many bunch crossings.

It is important to keep the LVL1 latency (time taken to form and distribute the LVL1 trigger decision) to a minimum. During this time, information for all detector channels has to be conserved in ‘pipeline’ memories. These memories are generally contained in custom integrated

circuits, placed on or close to the detector, usually in inaccessible regions and in a high-radiation environment. The total number of detector channels, excluding the pixel detectors, exceeds 10^7 . For reasons of cost and reliability, it is desirable to keep the pipeline lengths as short as possible. The LVL1 latency, measured from the time of the proton–proton collision until the trigger decision is available to the front-end electronics, is required to be less than $2.5 \mu\text{s}$. In order to achieve this, the LVL1 trigger is implemented as a system of purpose-built hardware processors. The target latency for the LVL1 trigger is $2.0 \mu\text{s}$ (leaving 500 ns contingency).

Events selected by LVL1 are read out from the front-end electronics systems of the detectors into readout drivers (RODs) and then into readout buffers (ROBs); present estimates foresee about 1700 ROBs in total. A large number of front-end electronics channels are multiplexed into each ROB. Intermediate buffers, labelled ‘derandomisers’ in Figure 1-2, average out the high instantaneous data rate at the output of the pipeline memories to match the available input bandwidth of the RODs.

All the detector data for the bunch crossing selected by the LVL1 trigger are held in the ROBs, either until the event is rejected by the level-2 (LVL2) trigger (in which case the data are discarded), or, in case the event is accepted by LVL2, until the data have been successfully transferred by the DAQ system to storage associated with the Event Filter (which makes the third level of event selection) [1-16]. The process of moving data from the ROBs to the Event Filter (EF) is called event building. Whereas before event building each event is composed of many fragments, with one fragment in each ROB, after event building the full event is stored in a single memory accessible by an EF processor.

The LVL2 trigger makes use of ‘region-of-interest’ (RoI) information provided by the LVL1 trigger. This includes information on the position (η and ϕ) and p_T of candidate objects (high- p_T muons, electrons/ γ , hadrons/ τ , jets), and energy sums (missing- E_T vector and scalar E_T value). The RoI data are sent by LVL1 to LVL2, for all events selected by the LVL1 trigger, using a dedicated data path. Using the RoI information, the LVL2 trigger selectively accesses data from the ROBs, moving only the data that are required in order to make the LVL2 decision. The LVL2 trigger has access to all of the event data, if necessary with the full precision and granularity. However, typically only data from a small fraction of the detector, corresponding to limited regions centred on the objects indicated by the LVL1 trigger, are needed by the LVL2 trigger. Hence, usually only a few per cent of the full event data are required thanks to this RoI mechanism.

It is expected that LVL2 will reduce the rate to $\sim 1 \text{ kHz}$. In contrast to the 75 kHz (upgradable to 100 kHz) limit for LVL1 that comes from the design of the detector front-end electronics, this is not a hard number. Optimisation of the sharing of the selection task between LVL2 and the EF is currently underway [1-17]. The latency of the LVL2 trigger is variable from event to event and is expected to be in the range $1\text{--}10 \text{ ms}$.

In the case of muon triggers, the rejection power at LVL2 comes from sharpening (and, where necessary, raising) the p_T threshold compared to LVL1, and from applying isolation requirements. Sharper p_T thresholds are obtained by using the precision muon chambers and the ID. The isolation requirements use the calorimeter information in a region around the muon candidate.

For isolated electrons, rejection power at LVL2 comes from using the full-granularity calorimeter information and requiring a matching high- p_T charged track in the ID; the transition-radiation signature provides additional rejection power. For photons, less rejection power is possible than in the case of electrons, since the ID cannot be used given the relatively high probability for

photon conversion in the ID material (it is not planned to use a track veto for the photon trigger). However, for the important physics channel $H \rightarrow \gamma\gamma$, the trigger can require a pair of photons, with a further rejection factor for each γ compared to LVL1 due to the use of high-precision, high-granularity calorimeter information.

For the hadron/ τ trigger, rejection at LVL2 is achieved using the full-granularity calorimeter information and the ID. A localised, isolated (hadronic) calorimeter cluster with a matching high- p_T track is required.

In the case of jets, much less rejection power is possible. Jets are the dominant high- p_T process at the LHC, and the threshold behaviour of the LVL1 trigger is reasonably sharp. Hence, for jet triggers, LVL2 must either increase the threshold, or make additional requirements in order to significantly reduce the rate. The possibility of identifying b-quark jets at LVL2 using the ID is under study.

Concerning the energy-sum triggers (E_T^{miss} , total scalar E_T), only limited improvement is possible using the RoI mechanism. The energy-sum values from LVL1 are provided to LVL2 and refinements can be made to correct, e.g. for high- p_T muons (the LVL1 E_T^{miss} trigger uses only calorimeter information, so muons contribute to the observed E_T^{miss}) or for saturated trigger-tower signals, a possibility allowed by design to reduce cost. The possibility of performing a full E_T^{miss} recalculation at LVL2 for a small subset of events remaining after other LVL2 selection criteria have already been applied is being investigated.

The LVL1 trigger makes available RoI information for all of the objects that contributed to the event being selected: these are called primary RoIs. In order to allow for additional requirements to be made at LVL2, the LVL1 trigger provides RoI information for objects that did not contribute to the selection of the event. Such RoIs, typically for objects of relatively low p_T , are called secondary RoIs.

After LVL2, the last stage of the online selection is performed by the EF. It will employ offline algorithms and methods, adapted to the online environment, and use the most up to date calibration and alignment information and the magnetic field map. The EF will make the final selection of physics events which will be written to mass storage for subsequent full offline analysis. The output rate from LVL2 should then be reduced by an order of magnitude, giving ~ 100 Hz, corresponding to an output data rate of ~ 100 MB/s if the full event data are to be recorded.

It is envisaged that the first task of the EF will be to confirm the results of the LVL2 decision and subsequently use the results of the LVL2 to seed its own analyses. The rejection power of the EF comes from:

- using refined algorithms and, where necessary, tighter p_T thresholds compared to those used in the LVL2;
- the availability of all data relevant to the specific event in calculations and selection criteria;
- the use of complex algorithms and criteria which, due to processing time limits, cannot be performed at LVL2, an example being vertex and track fitting using bremsstrahlung recovery for electrons.

1.8 Computing

Computing is crucial for the success of the ATLAS experiment and the effort to develop and maintain the software will be enormous [1-18]. The software is vitally important to the whole experiment's success, and must be maintained for a lifetime of about 20 years of the project. Hence, the quality requirements on the software have to be very high. It is estimated that, for the whole ATLAS software-development project, up to 1000 person-years will be required. Effort of such magnitude is only available in the collaboration on a very distributed basis. In order to optimise quality and to assure long-term maintenance, software development has to adhere to accepted international standards, seek common developments with other experiments, and employ commercial solutions wherever possible. It is planned to implement the software following the object-oriented paradigm. Currently, implementation using the C++ language is being pursued.

The anticipated data volume of about 1 PByte (10^{15} Bytes) per year requires new methods for data reduction, data selection, and data access for physics analysis. The basic goal is that every physicist in ATLAS must have the best possible access to the data to be analysed, irrespectively of his or her location. The proposed scheme consists of archiving the 'raw data' (1 PByte/year) selected by the Event Filter system. A first processing step will be performed on all data shortly (a few hours) after data taking. For this step, basic calibration and alignment constants must be available. The purpose is to evaluate basic physics quantities required by most analyses and to classify events into physics channels. The produced data have to be accessible at the event level and, below that, at the reconstruction and physics object level. ATLAS is considering using an object-oriented data-base system for this purpose. One copy of the data will be held at CERN, the replication of some or all of the data at a number of 'regional centres' is also foreseen.

To enable physics analysis to be performed in such a world-wide collaboration, high performance networking is a necessity. Today it is difficult to predict the evolution of the cost and performance of international networks at the time of LHC operation. As these are important parameters for the concrete planning of an analysis scenario, ATLAS is following the ongoing developments and will adjust its plans in these area accordingly.

1.9 References

- 1-1 'LHC White Book', CERN/AC/93-03; 'LHC Conceptual Design Report, CERN/AC/95-05.
- 1-2 ATLAS Collaboration, Technical Proposal for a General Purpose pp Experiment at the Large Hadron Collider at CERN, CERN/LHCC/94-43, LHCC/P2, 15 December 1994.
- 1-3 ATLAS Collaboration, Magnet System Technical Design Report, CERN/LHCC/97-18, 30 April 1997.
- 1-4 ATLAS Collaboration, Central Solenoid Technical Design Report, CERN/LHCC/97-21, 30 April 1997.
- 1-5 ATLAS Collaboration, End-Cap Toroids Technical Design Report, CERN/LHCC/97-20, 30 April 1997.
- 1-6 ATLAS Collaboration, Barrel Toroid Technical Design Report, CERN/LHCC/97-19, 30 April 1997.

- 1-7 ATLAS Collaboration, Inner Detector Technical Design Report, Volume 1, CERN/LHCC/97-16, 30 April 1997.
- 1-8 ATLAS Collaboration, Inner Detector Technical Design Report, Volume 2, CERN/LHCC/97-17, 30 April 1997.
- 1-9 ATLAS Collaboration, Pixel Detector Technical Design Report, CERN/LHCC/98-13, 31 May 1998.
- 1-10 ATLAS Collaboration, Calorimeter Performance Technical Design Report, CERN/LHCC/96-40, 15 December 1996.
- 1-11 ATLAS Collaboration, Liquid Argon Calorimeter Technical Design Report, CERN/LHCC/96-41, 15 December 1996.
- 1-12 ATLAS Collaboration, Tile Calorimeter Technical Design Report, CERN/LHCC/96-42, 15 December 1996.
- 1-13 ATLAS Collaboration, Muon Spectrometer Technical Design Report, CERN/LHCC/97-22, 31 May 1997.
- 1-14 ATLAS Collaboration, First-Level Trigger Technical Design Report, CERN/LHCC/98-14, 30 June 1998.
- 1-15 ATLAS Collaboration, Technical Coordination Technical Design Report, CERN/LHCC/99-01, 31 January 1999.
- 1-16 ATLAS Collaboration, DAQ, EF, LVL2 and DCS Technical Progress Report, CERN/LHCC/98-16, 30 June 1998.
- 1-17 ATLAS Collaboration, Trigger Performance Status Report, CERN/LHCC/98-15, 30 June 1998.
- 1-18 ATLAS Collaboration, Computing Technical Proposal, CERN/LHCC/96-43, 15 December 1996.